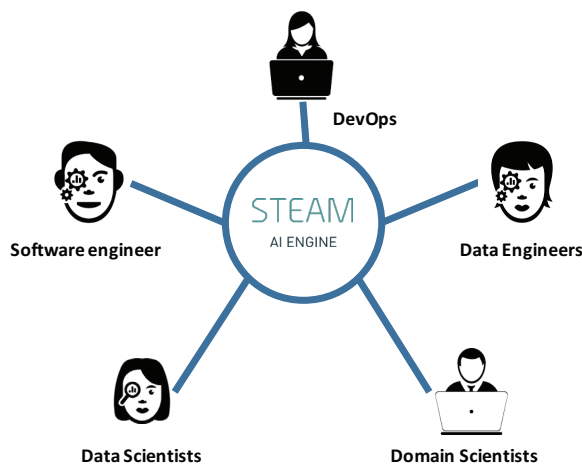


# Steam AI Engine

## What is Steam?

The Steam AI engine is an end-to-end platform that streamlines the entire process of building and deploying smart applications. Now data scientists and developers can launch turnkey compute environments for collaboratively training and deploying predictive models and integrate those models into real-time smart applications.



## Why Steam?

- Hard to collaborate between data scientists, software engineers and DevOps.
- Model comparison and collaboration is a pain
- Endless waiting for dedicated resources (e.g. red tape and security concerns) means further delay in getting business insights
- Updating models is an offline batch process and thus business cannot be agile in applying new models
- Most data science solutions are out of reach for teams without data science expertise

## Improve Business Efficiency

- Single platform for DevOps, data scientists, software engineers, and domain scientists to collaborate on.
- Support language of choice for different personas: R, Python, Java.
- Facilitate in-the-moment communication, reduce model deployment time and get to the results much faster.


## Improve Operational Resource Efficiency

- Shared infrastructure with multi-tenancy support
- ElasticML to elastically manage and change the size of underlying computing cluster
- Reduce your OPEX significantly

## A Platform to Build & Scale Smart Applications

- Elastic computing to allow data scientists to train on massive datasets on expandable clusters.
- Model management to share, collaborate and compare different model metrics.
- Streamline the entire process by providing an integrated hub that packages all underlying components into one easy-to-deploy piece of technology.


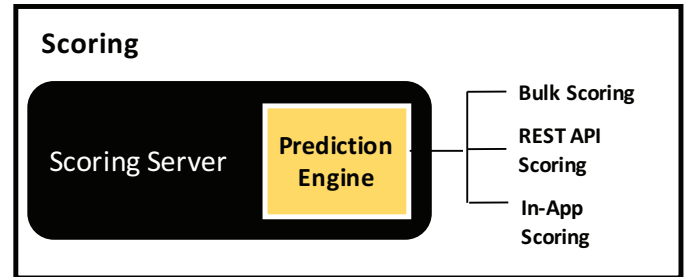




### Resource Sharing & Management

- Provide secure multi-tenant environments for users
- Automatically manage resources (e.g. memory) needed for the H2O instances without stopping or affecting existing tasks
- Auto scale-up/down of H2O instances and resources to achieve desired SLA for customers
- Start/stop cluster, allocate memory, start/pause/stop H2O instances

- Build and store thousands of models. Managed and promote them in Model Manager.
- Run A/B testing of Models
- Models can be built in R, Python and Java
- Easy specification of model management rules for model promotion either via GUI or programming language



- Bulk scoring of events in files increasing efficiency
- RESTful endpoints to operationalize scoring
- In-App Scoring for scoring within an app – js (app eng), Java (sys/Java app eng), python/R (data scientists)
- Deploy to real time streaming applications like Kafka, Spark Streaming etc.

To learn more contact [sales@h2o.ai](mailto:sales@h2o.ai)

### About H2O.ai

At H2O.ai we see a world where all software will incorporate AI, and we're focused on bringing AI to business through software. H2O.ai is the maker behind H2O, the leading open source machine learning platform for smarter applications and data products. H2O operationalizes data science by developing and deploying algorithms and models for R, Python and the Sparkling Water API for Spark. Some of H2O's mission critical applications include predictive maintenance, operational intelligence, security, fraud, auditing, churn, credit scoring, user based insurance, predicting sepsis, ICU monitoring and more in over 5,000 organizations. H2O is brewing a grassroots culture of data transformation in its customer communities. Customers include Capital One, Progressive Insurance, Zurich North America, Transamerica, Comcast, Nielsen Catalina Solutions, Neustar, Macy's, Walgreens, Kaiser Permanente and Aetna.

